
What we got wrong with IP

Jon Crowcroft

4.15pm, October 31, 2001

Computer Laboratory Seminar

Paraphrasing W. Churchill, one might say "The Internet - the worst network architecture, except for all the rest".

In this talk, I will look at some of the design mistakes (assuming that the Internet was actually designed, rather than discovered, like Tigger and Piglet finding the North Pole), that were made in the selection of algorithms and techniques in the lower levels of the Internet.

The talk will be structured around 20 key problems (well 23, in fact, for readers of the Illuminati trilogy), and will take in alternative solutions, and their relative merits and tradeoffs.

Specific areas of concern are pretty much everything, including naming, addressing, routing, reliability and security, as well as manageability and scaling.

At the end of the talk there will be time for questions, but probably not for answers.

“Twenty (sic.) Something Better”

From the movie “Roxanne,” starring Steve Martin.
based on a play, Cyrano de Bergerac by Edmund
Rostand...

The situation is that the fireman with a very large nose, who is extremely sensitive about it, has gone into a bar where everyone except the fall guy *know* that you *must* not even mention the nose (let alone make jokes about it. The foolish fall guy, for it is he, makes a lame attempt at a gag, at which point Steve Martin challenges him to a duel of jokes about the nose

Here, I’ve liberally adapted the jokes/insults, to act as a series of launch pads for attacks on the Internet Architecture (or “Internet Systems Design”, if you want a less pretentious name).

Obvious

1. Obvious: Excuse me. Is that your nose or did a bus park on your face.

The size of IP packet header is 40 bytes (TCP/IP or RTP/UDP/IP) - seems a bit poor when carrying ssh/telnet or GSM audio samples:-)

But we can always compress it - TCP and RTP header compression are not that complex, and can be customised to any typical situation - a nice employment opportunity for CS optimisation addicts.

Meteorological

2. Meteorological: Everybody take cover. She's going to blow.

Stability of routing- IP addresses are nearly global (pace, CIDR) and route computations are synchronised at each level (Inter-AS, and intra-AS) - so a local problem can brew a global storm.

In fact, its worse. BGP routes are asymmetric. This fits the "*business model*" of "getting the traffic *off of my net* as fast as possible. But its arbitrary.

Why can't we have any number of alternate business models? If routeing admitted of this, then we could avoid the appalling convergence time problems with BGP.*

*BGP takes $d * iBGP$ to converge, where $iBGP$ is the internal convergence time of an IGP e.g. 1 minute. d is typically the diameter of the Internet in ASs - around 14.

Fashionable

3. Fashionable: You know, you could de-emphasize your nose if you wore something larger. Like ... Wyoming.

MPLS is touted as a simplification (smaller *shim* headers, smaller labels to switch on. How come the gMPLS specification is 3 times the size of the IP one then?

IPv6 (DoCoMo) is trendy. Is it any better?

Well, Microsoft have an author on almost every IPv6 document. Cisco have one on every MPLS document. What does this tell us?

Personal

4. Personal: Well, here we are. Just the three of us.

Modulo the almost mythical IPSEC, wiretap appears not only easy, but almost criminally so.

But there's so much Internet measurement research that is now possible with tcpdump, ethereal etc etc

Punctual

5. Punctual: Alright gentlemen. Your nose was on time but you were fifteen minutes late.

The Internet offers what is amusingly known as the “best effort service”. Packets get there sometime or not. What time, or whether, is left to the users imagination.

But since there's no guarantees, there's no charge

Envious

6. Envious: Oooo, I wish I were you. Gosh. To be able to smell your own ear.

Reflective middleware (recursion in general - e.g. IP in IP in IP...) many possibilities for endless fun.

IP in IP tunnels are a form of overlays. In fact, IP was originally an overlay. Overlays are a low-entry-cost way to do novel network and distributed systems research. If you can do it on IP, you can do it anywhere!

Naughty

7. Naughty: Pardon me, Sir. Some of the ladies have asked if you wouldn't mind putting that thing away.

An oft cited report from CMU claims that 85% of content on the net is porn. Another oft cited report on AT&T's phone system is that 50% of their revenue is spent in collecting phone bills. Do we believe either of these?

If you can see it, you can catch the perps.

Philosophical

8. Philosophical: You know. It's not the size of a nose that's important. It's what's in it that matters.

Where IP may have a paucity of functionality, HTTP has an overabundance: WWW more than makes up for IP.

TCP research is almost endless (c.f. end-to-end arguments in systems design).

Humerous

9. Humorous: Laugh and the world laughs with you. Sneeze and its goodbye Seattle.

Location not discernible from IP address. More interestingly, DDOS attacks are really quite simple to deploy. Is this bad? Is being *slashdotted* bad? Are these questions different?

Anonymity is feasible.

Commercial

10. Commercial: Hi, I'm Earl Schibe and I can paint that nose for \$39.95.

Information is cheap.

Spam is expensive to the receiver. You get what you pay for.

Polite

11. Polite: Ah. Would you mind not bobbing your head. The orchestra keeps changing tempo.

The principle of unintended effects applies to everything.

Endless service creation opportunities including Professor McAuley's distributed orchestras.

Melodic

12. Melodic: Everybody! "He's got the whole world in his nose."

For some reason, whilst everyone claims that Cisco and other people implement FIFO, routers seem to cause packet replication even for single data grams with no link layer retransmission... ..

Software engineering ain't what it used to be:-)

Sympathetic

13. Sympathetic: Oh, What happened? Did your parents lose a bet with God?

Internet architecture considered gothic?

Well, anyone and everyone can play in the IETF, and do.

Complementary

14. Complementary: You must love the little birdies to give them this to perch on.

DOS attacks on IDS.

Coding speech in covert value signaling in TCP and IP spare fields

Scientific

15. Scientific: Say, does that thing there influence the tides.

what do IP packets weigh?

Forget the Interplanetary Internet - we already have enough black holes on the terrestrial superhighway.

Obscure

16. Obscure: Oh, I'd hate to see the grindstone.

TCP:tough on viruses: tough on the cause of viruses

Telemedicine, might eventually be useful...

Inquiry

17. Inquiry: When you stop to smell the flowers, are they afraid?

flow control? in a network? how very quaint...

Yes, there really is a MIME type for smell - application - distributed fish markets.

French

18. French: Say, the pigs have refused to find any more truffles until you leave.

IPSEC: pas devans les enfants

Well, actually they did repeal that part of the Code Napoleon.

Pornographic

19. Pornographic: Finally, a man who can satisfy two women at once.

multicast is often associated with promiscuity

but there are strict group management protocols too

Religious

20. Religious: The Lord giveth and He just kept on giving, didn't He.

DDOS

telemetry

Disgusting

21. Disgusting: Say, who mows your nose hair.

where do i give back an IP address when my machine breaks

with IPv6, you don't have to

Paranoid

22. Paranoid: Keep that guy away from my cocaine!
didn't Stajano work out a protocol for this...?

Aromatic

23. Aromatic: It must be wonderful to wake up in the morning and smell the coffee ... in Brazil.

MIME-TYPE: smell/coffee

Appreciative

24. Appreciative: Oooo, how original. Most people just have their teeth capped.

TV & telephones sort of work OK don't they

Dirty

25. Dirty: Your name wouldn't be Dick, would it?

I leave this to the audience's imagination.

Summary and Conclusions

the skeleton in the Internet closet, the dark horse...etc

The man with two brains

All of me